

PHASE TRANSITION OF THE CONTACT PROCESS ON RANDOM REGULAR GRAPHS

JEAN-CHRISTOPHE MOURRAT, DANIEL VALESIN

ABSTRACT. We consider the contact process with infection rate λ on a random $(d+1)$ -regular graph with n vertices, G_n . We study the extinction time τ_{G_n} (that is, the random amount of time until the infection disappears) as n is taken to infinity. We establish a phase transition depending on whether λ is smaller or larger than $\lambda_1(\mathbb{T}^d)$, the lower critical value for the contact process on the infinite, $(d+1)$ -regular tree: if $\lambda < \lambda_1(\mathbb{T}^d)$, τ_{G_n} grows logarithmically with n , while if $\lambda > \lambda_1(\mathbb{T}^d)$, it grows exponentially with n . This result differs from the situation where, instead of G_n , the contact process is considered on the d -ary tree of finite height, since in this case, the transition is known to happen instead at the *upper* critical value for the contact process on \mathbb{T}^d .

MSC 2010: 82C22; 05C80

KEYWORDS: INTERACTING PARTICLE SYSTEM, CONTACT PROCESS, RANDOM GRAPH, CONFIGURATION MODEL

1. INTRODUCTION

Let $G = (V, E)$ be a locally finite and connected graph. The contact process with parameter $\lambda > 0$ on a graph $G = (V, E)$ is a continuous-time Markov process on $\{0, 1\}^V$ with generator \mathcal{L} defined, for any local function $f : \{0, 1\}^V \rightarrow \mathbb{R}$ and any configuration $\xi \in \{0, 1\}^V$, by

$$(\mathcal{L}f)(\xi) = \sum_{x \in V} [f(\xi^{0 \rightarrow x}) - f(\xi)] + \lambda \sum_{x, y \in V} [\mathbf{1}_{\{\xi(x)=1\}} \cdot n(x, y) \cdot (f(\xi^{1 \rightarrow y}) - f(\xi))],$$

where

$$\xi^{i \rightarrow x}(z) = \begin{cases} i & \text{if } z = x; \\ \xi(z) & \text{otherwise,} \end{cases}$$

$\mathbf{1}$ is the indicator function, and $n(x, y)$ is the number of edges between x and y (which we allow to be larger than 1).

Vertices of the graph are interpreted as individuals and states 0 and 1 indicate that an individual is healthy or infected, respectively. The transitions that appear in the definition of \mathcal{L} are then understood as follows: $\xi \rightarrow \xi^{0 \rightarrow x}$ means that there is a *recovery* at x , and $\xi \rightarrow \xi^{1 \rightarrow x}$ means that there is a *transmission* to x (this is necessarily triggered by a neighbour $y \sim x$ with $\xi(y) = 1$). We refer the reader to [Li99] for a thorough introduction to the contact process, including the facts that we state without proof in this introduction.

We often abuse notation and, for a set S , treat an element $\xi \in \{0, 1\}^S$ as the same as the set $\{x \in S : \xi(x) = 1\}$; in particular, the configuration identically equal to zero is denoted by \emptyset . We denote by $(\xi_t^A)_{t \geq 0}$ the contact process on G with initial configuration $\xi_0^A \equiv A$. If $A = \{x\}$, we write (ξ_t^x) instead of $(\xi_t^{\{x\}})$. We write $P_{G, \lambda}$ for a probability measure under which the contact process with parameter λ on G is defined.

The *extinction time* τ_G^A for the contact process on $G = (V, E)$ with initial configuration $A \subseteq V$ is defined by

$$\tau_G^A = \inf\{t : \xi_t^A = \emptyset\}.$$

Since the dynamics only allows for new infections to appear by transmission, the configuration \emptyset is absorbing, so we have $\xi_t^A = \emptyset$ for any $t \geq \tau_G^A$. We write τ_G instead of τ_G^G and τ_G^x instead of $\tau_G^{\{x\}}$.

When the contact process is considered on infinite graphs, a central question is whether we have survival or extinction of the infection. For a graph $G = (V, E)$, a finite set $A \subseteq V$ and $\lambda > 0$, define

$$\begin{aligned} p_{G,A,\lambda}^{\text{ext}} &= P_{G,\lambda}[\tau_G^A < \infty] = P_{G,\lambda}[\exists t_0 : \xi_t^A = \emptyset \text{ for all } t \geq t_0], \\ p_{G,A,\lambda}^{\text{loc ext}} &= P_{G,\lambda}[\exists t_0 : \xi_t^A \cap A = \emptyset \text{ for all } t \geq t_0], \end{aligned}$$

the probabilities of *extinction* and *local extinction* of the contact process with parameter λ on G started from $\xi_0^A \equiv A$. It can be shown that we either have $p_{G,A,\lambda}^{\text{ext}} = 1$ for all A or $p_{G,A,\lambda}^{\text{ext}} < 1$ for all A ; likewise, either $p_{G,A,\lambda}^{\text{loc ext}} = 1$ for all A or $p_{G,A,\lambda}^{\text{loc ext}} < 1$ for all A . We say that the contact process

- *dies out* if $p_{G,A,\lambda}^{\text{ext}} = 1$ for all A ;
- *survives weakly* (or *globally but not locally*) if $p_{G,A,\lambda}^{\text{ext}} < 1$ and $p_{G,A,\lambda}^{\text{loc ext}} = 1$ for all A ;
- *survives strongly* (or *locally*) if $p_{G,A,\lambda}^{\text{loc ext}} < 1$ for all A .

We let $\lambda_1(G) = \sup\{\lambda : p_{G,A,\lambda}^{\text{ext}} = 1\}$ and $\lambda_2(G) = \sup\{\lambda : p_{G,A,\lambda}^{\text{loc ext}} = 1\}$. It is well known that for $G = \mathbb{Z}^d$, $0 < \lambda_1(\mathbb{Z}^d) = \lambda_2(\mathbb{Z}^d) =: \lambda_c(\mathbb{Z}^d) < \infty$, whereas for $G = \mathbb{T}^d$, the infinite, $(d+1)$ -regular tree (with $d \geq 2$), the situation is quite different, as we then have $0 < \lambda_1(\mathbb{T}^d) < \lambda_2(\mathbb{T}^d) < \infty$. In the latter case, if we take λ in the weak survival regime and start the process from a single infection at a vertex x , the infection has a chance of surviving, but it can only do so by propagating outwards from x ; any finite neighbourhood of x only carries the infection for a finite amount of time, as required in the definition of weak survival.

The contact process on finite graphs (deterministic or random) has also been the subject of much investigation; below we will survey some of the past work that is most relevant to the object of interest of this paper. Let us start noting that for finite graphs, there is no question of survival or extinction: if G is finite, the infection almost surely disappears in finite time. One is thus interested in the behaviour of the infection before it disappears. The typical course of action goes as follows: we fix λ and the parameters that define the graph, then take a size parameter (e.g. the number of vertices) to infinity, so as to obtain a sequence of graphs G_n , and consider the asymptotic behaviour of some random quantity X_n associated with the contact process on G_n . Common choices for X_n are the extinction time τ_{G_n} and the average proportion of infected vertices before τ_{G_n} . Of particular interest are the cases in which the asymptotic behaviour of X_n depends sensitively on the choice of λ and the set of parameters that define the graph; this dependency can sometimes be related to a phase transition of the contact process on an infinite graph that is in some sense approximated by G_n .

This project has been first carried out for G_n equal to the subgraph of \mathbb{Z}^d induced by the vertices contained in a box of side length n ; see [CGOV84, Sc85, DL88, Ch94, DS88, Mo93, Mo99], or [Li99, Section I.3] for an overview. As suggested in the previous paragraph, the behaviour of τ_{G_n} exhibits a phase transition that mimics the phase transition of the contact process on \mathbb{Z}^d : τ_{G_n} grows logarithmically or exponentially with the volume of G_n respectively if λ is taken smaller or larger than $\lambda_c(\mathbb{Z}^d)$.

Let us now turn to finite trees. Let \mathbb{T}_ℓ^d denote the d -ary tree of height ℓ . In other words, \mathbb{T}_ℓ^d is a tree with a distinguished vertex o (called the root) so that o has degree d , all vertices at distance between 1 and $\ell - 1$ from o have degree $d + 1$ and all vertices at distance ℓ from o have degree 1. In [St01] and [CMMV13], the contact process on \mathbb{T}_ℓ^d was studied. In [St01], bounds were obtained on $\tau_{\mathbb{T}_\ell^d}$ for different values of λ , and these bounds were improved in [CMMV13] to yield the result below. We denote by $|A|$ the cardinality of a set A .

Theorem 1.1. [CMMV13] *Assume $d \geq 3$.*

- (a) *For any $\lambda \in (0, \lambda_2(\mathbb{T}^d))$, there exists $c \in (0, \infty)$ such that, as $\ell \rightarrow \infty$,*

$$\frac{\tau_{\mathbb{T}_\ell^d}}{\log |\mathbb{T}_\ell^d|} \rightarrow c \text{ in probability.}$$

- (b) *For any $\lambda \in (\lambda_2(\mathbb{T}^d), \infty)$, there exists $C \in (0, \infty)$ such that, as $\ell \rightarrow \infty$,*

$$\frac{\log \tau_{\mathbb{T}_\ell^d}}{|\mathbb{T}_\ell^d|} \rightarrow C \text{ in probability.}$$

Moreover, $\tau_{\mathbb{T}_\ell^d}$ divided by its expectation converges in distribution to the exponential distribution with parameter 1.

In short, the asymptotic behaviour of the extinction time on \mathbb{T}_ℓ^d has a phase transition at the value $\lambda = \lambda_2(\mathbb{T}^d)$. It is worth mentioning that the above theorem remains true if \mathbb{T}_ℓ^d is replaced by $\hat{\mathbb{T}}_\ell^d$, defined as the subgraph of \mathbb{T}_ℓ^d induced by the vertex set given by a distinguished root vertex o and all vertices at graph distance at most ℓ from o (\mathbb{T}_ℓ^d is obtained from $\hat{\mathbb{T}}_\ell^d$ by removing one of the sub-trees emanating from the root of \mathbb{T}_ℓ^d).

The main goal of this paper is to investigate the asymptotic behaviour of the extinction time of the contact process on random $(d + 1)$ -regular graphs, which we now describe. Fix $n \in \mathbb{N}$, the number of vertices, with the restriction that $n(d + 1)$ be even. Let us define a random graph G_n with vertex set $V_n = \{1, \dots, n\}$. We endow each vertex with $d + 1$ half-edges; attaching two half-edges (say, one belonging to vertex i and another to vertex j) produces an edge between i and j . Then, any perfect matching on the set of half-edges produces a graph. We choose one perfect matching uniformly at random from all the $(n(d + 1) - 1) \cdot (n(d + 1) - 3) \cdots 3 \cdot 1$ possibilities, and this produces the random edge set E_n . Obviously, G_n is a $(d + 1)$ -regular graph; by using an alternate construction of G_n that is described in Section 2, it is easy to show that, for any R , with probability tending to 1 as $n \rightarrow \infty$, the ball of radius R around a given vertex has no loop, and is thus isomorphic to $\hat{\mathbb{T}}_\ell^d$.

We write \mathbb{P}_n for the law of G_n , and $\mathbb{P}_{n,\lambda}$ for a probability measure under which both G_n and the contact process on G_n with parameter λ are defined. Our main result is

Theorem 1.2. *Assume $d \geq 3$.*

- (a) *For every $\lambda \in (0, \lambda_1(\mathbb{T}^d))$ there exists $C < \infty$ such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}_{n,\lambda} [\tau_{G_n} < C \log(n)] = 1.$$

- (b) *For every $\lambda \in (\lambda_1(\mathbb{T}^d), \infty)$ there exists $c > 0$ such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}_{n,\lambda} [\tau_{G_n} > e^{cn}] = 1.$$

Let us stress the most important point: contrary to the situation in Theorem 1.1, the behaviour of the extinction time has here a phase transition at $\lambda_1(\mathbb{T}^d)$ instead of $\lambda_2(\mathbb{T}^d)$. Although surprising at first, this phenomenon can be understood as follows: when ℓ is large, the tree \mathbb{T}_ℓ^d as seen from a vertex chosen uniformly at random does

not at all locally look like \mathbb{T}^d . For example, with non-vanishing probability, the random vertex is a leaf of \mathbb{T}_ℓ^d . The notion of local limit of a sequence of graphs, as formalized in [BS01], captures such features, and one can check that \mathbb{T}_ℓ^d does not converge locally to \mathbb{T}^d , but rather to another infinite graph called the *canopy tree* \mathbb{CT}^d (see [Be, Example 5.14] for details). Moreover, one can show that

$$(1.1) \quad \lambda_1(\mathbb{CT}^d) = \lambda_2(\mathbb{CT}^d) = \lambda_2(\mathbb{T}^d)$$

(local survival for $\lambda > \lambda_2(\mathbb{T}^d)$ can be derived from [CMMV13, Theorem 1.6], while extinction for $\lambda \leq \lambda_2(\mathbb{T}^d)$ follows from [Li99, Proposition 4.57]). In light of this, the logarithmic-exponential phase transition for the extinction time on \mathbb{T}_ℓ^d does happen at the *lower* critical value of the limiting graph after all, but one must take into account that this limiting graph is the canopy tree.

It would be very interesting to find wider classes of graphs (such as the configuration model described below) for which one can prove that the same phase transition occurs at the lower critical value of the limiting graph. This is reminiscent of the question of the locality of the percolation critical probability, see [BNP11], [Be, Section 5.2] and [MT13].

The graph G_n is a particular case of the class of graphs known as the *configuration model*. Whereas for G_n we assumed that all degrees are taken equal to $d + 1$, in the configuration model one allows for a random degree sequence, typically i.i.d. from some fixed degree distribution. The contact process on these graphs has lately received attention ([CD09, MMVY12, MVY13]). In particular, it was proved that, if the degree distribution is a power law (and the graph is assumed to be connected), then the extinction time grows exponentially with the number of vertices *regardless of the value of λ* . In this sense, one can say that the contact process on the configuration model with a power law degree distribution has no phase transition: it is “always supercritical”. Theorem 1.2 shows that, if the degrees are constant, then there is a phase transition. In fact, putting together Lemma 3.2 below with the main result of [MMVY12], we get

Theorem 1.3. *Let $\mathcal{G}_{n,d}$ be the set of connected graphs with n vertices and degree bounded by $d + 1$.*

(a) *For any $\lambda \in (0, \lambda_1(\mathbb{T}^d))$, there exists $C < \infty$ such that*

$$\lim_{n \rightarrow \infty} \inf_{G \in \mathcal{G}_{n,d}} P_{G,\lambda} [\tau_G < C \log n] = 1.$$

(b) *For any $\lambda \in (\lambda_c(\mathbb{Z}), \infty)$, there exists $c > 0$ such that*

$$\lim_{n \rightarrow \infty} \inf_{G \in \mathcal{G}_{n,d}} P_{G,\lambda} [\tau_G > e^{cn}] = 1.$$

In particular, if we take the contact process on the configuration model with a degree distribution with *bounded support*, the extinction time exhibits a phase transition on λ . An open problem that we believe to be of much interest is whether such a phase transition also occurs if the degree distribution has unbounded support, but a light tail. Equally interesting (and difficult) is the corresponding question for supercritical Erdős-Rényi random graphs.

We now comment on the proofs of the two parts of Theorem 1.2. Part (a) relies on the fact, stated in Lemma 3.2, that the contact process on any graph of degree bounded by $d + 1$ is stochastically dominated (in terms of the number of infected vertices) by a contact process with same λ on \mathbb{T}^d . Although this is a very elementary and natural statement, we did not find it in the literature, so we give a proof, which relies on the concept of universal covering of a graph. Once Lemma 3.2 has been established, we invoke known bounds for the extinction time of the contact process on \mathbb{T}^d in the extinction regime in order to conclude the proof.

As a side note, we remark that Lemma 3.2 also applies to infinite graphs, and as a consequence we prove:

Theorem 1.4. *If $\lambda \leq \lambda_1(\mathbb{T}^d)$, then the contact process with parameter λ on any graph with degree bounded by d dies out.*

The proof of part (b) of Theorem 1.2 is much more involved, and relies on a geometric property of G_n (Theorem 2.4). The idea can be roughly summarised as follows. Assume that at time 0, there are εn infected vertices ($\varepsilon > 0$ small) with the property that within distance r of each of these vertices, one can find a copy of \mathbb{T}_ℓ^d , the copies (and the paths to them) being pairwise disjoint. Let us say that a set satisfying this property is *regenerative*. If ℓ can be chosen sufficiently large, then after some time, the infection will spread to $k\varepsilon n$ vertices with k large, except for an event whose probability is exponentially small in n . The crucial geometric property of G_n that we need is that if k is sufficiently large, then with probability tending to 1, any subset of size $k\varepsilon n$ contains a regenerative subset of size εn . This enables to iterate the argument. Since the probability of failure of a given step is exponentially small, the extinction time must be exponentially long.

In the rest of this introduction, we give a summary of the notation we use and an exposition of the graphical construction of the contact process. Sections 2 and 3 are devoted to proving parts (b) and (a) of Theorem 1.2 respectively.

1.1. Summary of notation for sets and graphs. For a set A , we write $|A|$ for the cardinality of A and $\mathbf{1}_A$ for the indicator function of A .

A *graph* is a pair $G = (V, E)$, where V is the set of *vertices* and E is the set of *edges*. With each edge is associated an unordered pair of points in V , the end-points of the edge. Loops (i.e. edges between a vertex and itself) and multiple edges with the same end-points are allowed. We may abuse notation and write $v \in G$ instead of $v \in V$ when convenient. The *degree* of a vertex is the number of non-loop edges that contain it plus twice the number of loops that contain it. The *graph distance* between vertices x and y is the length of a shortest path between x and y , and is denoted by $\text{dist}(x, y)$, or $\text{dist}_G(x, y)$ when we want to make the graph explicit.

The set \vec{E} of oriented edges of G is defined as follows: for each $e \in E$, we add to \vec{E} two oriented edges corresponding to the two possible orientations of e . A generic element of \vec{E} will be denoted by \vec{e} . Note that, if there are k edges containing x and y , then there are $2k$ oriented edges containing x and y . We write $v_0(\vec{e})$ and $v_1(\vec{e})$ to denote the starting and ending vertex of \vec{e} , respectively. We also let $u(\vec{e}) \in E$ denote the unoriented edge to which \vec{e} is associated.

A *rooted graph* is a pair (ρ, G) where G is a graph and $\rho \in G$. Given two rooted graphs (ρ, G) and (ρ', G') with $G = (V, E)$, $G' = (V', E')$, we say that $f : V \rightarrow V'$ is an *embedding* of (ρ, G) into (ρ', G') if

- $f(\rho) = \rho'$,
- f is injective,
- for every $u, v \in V$, u and v are neighbours in G if and only if $f(u)$ and $f(v)$ are neighbours in G' .

We say that (ρ', G') *embeds* (ρ, G) if there exists an embedding of (ρ, G) into (ρ', G') ; that (ρ, G) and (ρ', G') are *isomorphic* if there exists a bijective embedding from one to the other.

For the rest of the paper, we fix $d \geq 3$ and omit d from \mathbb{T}^d and \mathbb{T}_ℓ^d , thus writing \mathbb{T} and \mathbb{T}_ℓ respectively. The root of \mathbb{T}_ℓ is denoted by o .

1.2. Graphical construction of the contact process. Almost every paper on the contact process contains a brief exposition of its graphical construction, so this is by now quite a redundant addition. We nevertheless include it here because we

allow our graphs to contain loops and parallel edges, so we need to be careful with the notation to avoid confusion.

Fix a graph $G = (V, E)$ and an infection rate $\lambda > 0$. Assume given the following families of independent Poisson point processes on $[0, \infty)$: $(T_{\vec{e}})_{\vec{e} \in \vec{E}}$, all with rate λ , and $(D_x)_{x \in V}$, all with rate 1. We now imagine G is embedded on the xy -plane and add marks on the vertical lines $(V \cup \vec{E}) \times [0, \infty)$ as follows: for each x and each $t \in D_x$, we add a so-called *recovery mark* at (x, t) , and for each \vec{e} and each $t \in T_{\vec{e}}$, we add a so-called *transmission arrow* at (\vec{e}, t) .

Given $x, y \in V$ and $0 \leq s \leq t$, we say that (x, s) and (y, t) are connected by an *infection path* (and write $(x, s) \leftrightarrow (y, t)$) if there exists a right-continuous function $\gamma : [s, t] \rightarrow V$ such that

- $\gamma(s) = x, \gamma(t) = y$,
- $(r, \gamma(r)) \notin D_{\gamma(r)}$ for all $r \in [s, t]$,
- $\gamma(r-) \neq \gamma(r)$ implies $r \in T_{\vec{e}}$ for some \vec{e} with $v_0(\vec{e}) = \gamma(r-)$, $v_1(\vec{e}) = \gamma(r)$.

In other words, $(r, \gamma(r))_{s \leq r \leq t}$ must be a path from (x, s) to (y, t) which does not cross any recovery mark and is allowed to traverse transmission arrows.

For $x \in V$, let $\xi_t^x = \{y \in V : (x, 0) \leftrightarrow (y, t)\}$; for $A \subseteq V$, let $\xi_t^A = \cup_{x \in A} \xi_t^x$. Then, $(\xi_t^A)_{t \geq 0}$ is a Markov process on $\{0, 1\}^V$ with the same distribution as the one given by the generator \mathcal{L} defined earlier. The graphical construction has the advantage that all the contact processes $((\xi_t^A)_{t \geq 0})_{A \subseteq V}$ are defined in the same probability space and, if $A \subseteq B$, then $\xi_t^A \subseteq \xi_t^B$ for all t . Note also that $\xi_{s+t}^A = \cup_{x \in \xi_t^A} \{y : (x, t) \leftrightarrow (y, t+s)\}$.

2. SUPERCRITICAL REGIME

The goal of this section is to prove part (b) of Theorem 1.2. We first state for later usage a classical large deviation result on binomial random variables, whose proof is standard.

Lemma 2.1. *Let $\text{Bin}(m, p)$ denote a binomial random variable with parameters $m \in \mathbb{N}$ and $p \in [0, 1]$ (defined with respect to the measure \mathbb{P} , say). For every $\delta \geq 0$,*

$$\mathbb{P}[\text{Bin}(m, p) \geq (p + \delta)m] \leq e^{-m\psi_p(\delta)},$$

where

$$\begin{aligned} \psi_p(\delta) &= \sup_{\lambda} [\lambda(p + \delta) - \log(1 - p + pe^\lambda)] \\ (2.1) \quad &= (p + \delta) \log\left(\frac{p + \delta}{p}\right) + (1 - p - \delta) \log\left(\frac{1 - p - \delta}{1 - p}\right). \end{aligned}$$

We now recall the following important estimate from [SS98].

Lemma 2.2. [SS98] *For every $\lambda > \lambda_1(\mathbb{T})$, there exists $S, p_0 > 0$ and $\alpha > 1$ such that for every ℓ large enough,*

$$P_{\mathbb{T}, \lambda} [|\{x \in \xi_{\ell S}^o : \text{dist}(o, x) = \ell\}| \geq \alpha^\ell] \geq p_0,$$

We will need a slight modification of the above result:

Lemma 2.3. *For every $\lambda > \lambda_1(\mathbb{T})$ and $r > 0$, there exist $R, \sigma > 0$ and $\alpha > 1$ such that for every ℓ large enough, the following holds. For any graph G with vertices x, y so that $\text{dist}(x, y) \leq r$ and (y, G) embeds (o, \mathbb{T}_ℓ) ,*

$$P_{G, \lambda} [|\xi_{R\ell}^{\{x\}}| \geq \alpha^\ell] > \sigma.$$

Proof. This follows from Lemma 2.2 and the facts that

$$P_{G, \lambda} [\xi_t^{\{x\}}(x) = 1] \geq e^{-t}, \quad P_{G, \lambda} [\xi_{\text{dist}(x, y)}^{\{x\}}(y) = 1] \geq (e^{-2}(1 - e^{-\lambda}))^{\text{dist}(x, y)}.$$

The first estimate is obtained by considering the event that x does not recover, by time t , of the infection it has at time 0. The second estimate is obtained by considering the event that, in every unit time interval $[k, k+1]$ from time 0 to time $t = \text{dist}(x, y)$, the infection advances one step along a shortest path from x to y (given a pair of vertices u, v such that u is infected at time k , v is sure to be infected at time $k+1$ if, in this unit time interval, neither of them recover and there is at least one transmission from u to v). \square

We say that a set of vertices $W \subseteq V_n$ is (ℓ, r) -regenerative if there exists a family $(G'_v)_{v \in W}$ of subgraphs of G_n that are pairwise disjoint and such that for every $v \in W$, the following two conditions hold:

- G'_v contains v ,
- there exists $x \in G'_v$ such that the distance in G'_v between x and v is r and (x, G'_v) embeds (o, \mathbb{T}_ℓ) .

The crucial geometric property of G_n that we need is

Theorem 2.4 (Finding large regenerative subsets). *For k and r sufficiently large and for every ℓ , there exists $\varepsilon_0 > 0$ such that for every $\varepsilon \leq \varepsilon_0$, the following holds with \mathbb{P}_n -probability tending to 1 as n tends to infinity. From every $W \subseteq V_n$ of cardinality at least $k\varepsilon n$, one can extract an (ℓ, r) -regenerative subset of cardinality at least εn .*

Remark 2.5. In what follows, in order to prevent the notation from getting too heavy, we will pretend that certain quantities, such as εn and $k\varepsilon n$ in the above proposition, are integers. It should be clear that in a correct but overscrupulous writing, one should take the pertinent integer parts or add or subtract 1 at certain places, but that this does not change our proofs in any relevant way.

We prove part (b) of Theorem 1.2 using Theorem 2.4, before turning to the proof of the latter.

Proof of part (b) of Theorem 1.2. We fix $\lambda > \lambda_1(\mathbb{T})$, and choose constants according to the following steps:

- (1) fix r and k large, as required by Theorem 2.4;
- (2) let α, R, σ correspond to λ and r , as in Lemma 2.3;
- (3) take ℓ large enough, as required by Lemma 2.3, and also so that $\alpha^\ell > \frac{2k}{\sigma}$;
- (4) take $\varepsilon \leq \varepsilon_0$, where ε_0 corresponds to k, r, ℓ as in Theorem 2.4.

Assume G_n satisfies the property stated in Theorem 2.4, namely

every $W \subseteq V_n$ with $|W| \geq k\varepsilon n$ has an (ℓ, r) -regenerative subset of cardinality εn .

We will now prove that, for some constant $c > 0$ which does not depend on n ,

$$(2.2) \quad \text{for all } W \subseteq V_n \text{ with } |W| \geq k\varepsilon n, \quad P_{G_n} [|\xi_{R\ell}^W| \geq k\varepsilon n] \geq 1 - e^{-cn}.$$

This will imply the statement of Theorem 1.2.

We fix W with $|W| \geq k\varepsilon n$ and extract from it an (ℓ, r) -regenerative subset W' of cardinality εn ; we enumerate its elements, $W' = \{v_1, \dots, v_{\varepsilon n}\}$. By the definition of (ℓ, r) -regenerative sets, there exist pairwise disjoint subgraphs of G , $G'_{v_1}, \dots, G'_{v_{\varepsilon n}}$ such that $v_i \in G'_{v_i}$ and there exists $x_i \in G'_{v_i}$ such that $\text{dist}(v_i, x_i) \leq r$ and (x_i, G'_{v_i}) embeds (o, \mathbb{T}_ℓ) .

For each v_i , let $(\zeta_t^{v_i})_{t \geq 0}$ be the contact process on G'_{v_i} , started from only v_i infected, and built using the same family of Poisson processes as the original contact process on G . If $i \neq j$, then $(\zeta_t^{v_i})$ and $(\zeta_t^{v_j})$ are independent (since the G'_{v_i} are disjoint) and moreover, $\xi_t^W \supseteq \xi_t^{W'} \supseteq \cup_{v \in W'} \zeta_t^v$ for all t .

Define the events

$$E_i = \{|\zeta_{R\ell}^{v_i}| \geq \alpha^\ell\}, \quad 1 \leq i \leq \varepsilon n.$$

We then have $P_{G_n}[E_i] \geq \sigma$. Thus, by standard large deviation estimates for binomial random variables (see Lemma 2.1), we have

$$P_{G_n} \left[\sum_{i=1}^{\varepsilon n} \mathbf{1}_{E_i} \geq \frac{\sigma}{2} \varepsilon n \right] \geq 1 - e^{-c(\varepsilon, \sigma)n}.$$

Finally, if the event in the above probability occurs, we have

$$|\xi_{R\ell}^W| \geq \alpha^\ell \cdot \sum_{i=1}^{\varepsilon n} \mathbf{1}_{E_i} \geq \alpha^\ell \cdot \frac{\sigma}{2} \varepsilon n > k\varepsilon n.$$

This completes the proof. \square

The proof of Theorem 2.4 goes by showing first a seemingly weaker property, where the embedded trees are allowed to be somewhat damaged. We now introduce the relevant notions.

We call *pruned ℓ -tree* any graph obtained by removing an edge from \mathbb{T}_ℓ and taking the connected component of the root. We call *rooted pruned ℓ -tree* any rooted graph (o, H) , where H is a pruned tree obtained from \mathbb{T}_ℓ and o is the root of \mathbb{T}_ℓ (which belongs to H by construction).

We say that a rooted graph (ρ, H) is (ℓ, r) -favourable if there exists $x \in H$ such that the following two conditions hold:

- the distance between x and ρ is r ,
- the rooted graph (x, H) embeds a rooted pruned ℓ -tree.

We say that a set of vertices $W \subseteq V_n$ is (ℓ, r) -good if there exists a family $(G'_v)_{v \in W}$ of subgraphs of G_n whose vertex sets are pairwise disjoint such that for every $v \in W$, G'_v contains v and (v, G'_v) is (ℓ, r) -favourable.

Proposition 2.6 (Finding large good subsets). *For k and r sufficiently large and for every ℓ , there exists $\varepsilon_0 > 0$ such that for every $\varepsilon \leq \varepsilon_0$, the following holds with \mathbb{P}_n -probability tending to 1 as n tends to infinity. From every $W \subseteq V_n$ of cardinality $k\varepsilon n$, one can extract an (ℓ, r) -good subset of cardinality εn .*

We say that a set of vertices $W \subseteq V$ is r -prepared if the following two properties hold:

- for every $v \in W$, the r -neighbourhood of v is loop-free (we call r -neighbourhood of a site $v \in V$ the set of sites w whose distance to v is at most r),
- for every two distinct $v, w \in W$, the r -neighbourhoods of v and w are disjoint.

The key step of the proof lies in the following proposition.

Proposition 2.7 (Key estimate). *For every $k > 1$ and $\eta < 2 - 4/k$, there exists $r_0 \in \mathbb{N}$ such that the following holds. For every $r \geq r_0$ and $\ell \in \mathbb{N}$, there exists $c > 0$ and $\varepsilon_0 > 0$ such that for every $\varepsilon \leq \varepsilon_0$ and every n sufficiently large, the \mathbb{P}_n -probability that the set $W = \{1, \dots, k\varepsilon n\} \subseteq V_n$ is r -prepared but does not contain any (ℓ, r) -good subset of size εn is smaller than*

$$(c\varepsilon^\eta)^{|W|}.$$

Before proving this key estimate, we introduce some terminology. A *semi-graph* $g = (V, \mathcal{E}, \mathcal{H})$ is a triple consisting of a set V of vertices, a set \mathcal{E} of edges between points of V , and a set \mathcal{H} of *half-edges*, each half-edge being attached to some vertex in V . Given two half-edges h and h' , we write $h + h'$ to denote the edge obtained by “gluing together” the half-edges h and h' (that is to say, if h is attached to a vertex v and h' to a vertex w , then $h + h'$ is an edge connecting v and w). The distance in the semi-graph g is simply the distance in the graph (V, \mathcal{E}) .

Consider the semi-graph $g = (V_n, \mathcal{E}, \mathcal{H})$ such that $\mathcal{E} = \emptyset$ and each vertex in V_n has exactly $d + 1$ half-edges. We can construct a random regular graph with distribution \mathbb{P}_n by the following recursive procedure. Take an arbitrary half-edge h in \mathcal{H} (call it the *elected* half-edge); take a half-edge h' uniformly at random in $\mathcal{H} \setminus \{h\}$; add the edge $h + h'$ to the set \mathcal{E} , remove the half-edges h and h' from \mathcal{H} ; repeat until the set of half-edges is empty. (Recall that we assume $(d + 1)n$ to be even, so that there is an even number of half-edges to begin with.) The resulting random graph has law \mathbb{P}_n . A feature of this procedure that will be crucial in our reasoning is that at each step, we have the freedom to choose the elected half-edge as we wish among the half-edges of \mathcal{H} .

It will be convenient to write operations on sets such as those done in the above construction in the more symbolic form

$$(2.3) \quad \mathcal{E} \leftarrow \mathcal{E} \cup \{h + h'\}, \quad \mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\}.$$

Proof of Proposition 2.7. We take the semi-graph $g = (V_n, \mathcal{E}, \mathcal{H})$ such that $\mathcal{E} = \emptyset$ and such that every site has exactly $(d + 1)$ half-edges. As long as there is a half-edge h attached to a site at distance strictly less than r from \mathcal{W} , we choose this as the elected edge, pick another edge h' uniformly at random from $\mathcal{H} \setminus \{h\}$, and do the operations in (2.3). When there is no longer any such half-edge, the graph is still incomplete, but we can already decide if the set \mathcal{W} is r -prepared, since we have constructed the r -neighbourhood of \mathcal{W} (and nothing more). If the set is not r -prepared, we can stop. On the event that the set is r -prepared, we continue our construction of the graph, with the aim to show that \mathcal{W} will contain a good subset of size εn with high probability.

Let \mathcal{F} be the set of sites still having $(d + 1)$ half-edges at this point. We call elements of \mathcal{F} *fresh* vertices (they are still unseen by our construction of the graph). Let us call an element of \mathcal{W} a *seed*, and a site at distance r from a seed a *bud*. Recall that since we consider only the event that \mathcal{W} is r -prepared, to every seed corresponds exactly $(d + 1)d^{r-1}$ buds (so in total we have $|\mathcal{W}|(d + 1)d^{r-1}$ buds). We say that a bud is *active* if it has d half-edges. We say that a seed is *active* if it is associated with an active bud (i.e. if it is at distance r from an active bud). A seed or a bud that is not active will be called *quiet*. As of now, every bud is active, and a fortiori every seed is active. We let ρ be any active seed (as of now, this only means that we take $\rho \in \mathcal{W}$), and run the Pass as described below.

The Pass started from the active seed ρ .

By definition of being an active seed, ρ is associated with at least one active bud, say x , whose set of d attached half-edges we write $\overline{\mathcal{H}}$. We let $\overline{V} = \{x\}$, $\overline{\mathcal{E}} = \emptyset$.

Step 1

Let $v \in \overline{V}$ be at distance strictly less than ℓ from x as measured in the graph $(\overline{V}, \overline{\mathcal{E}})$, and such that there is a half-edge $h \in \overline{\mathcal{H}}$ attached to it. If such a v does not exist, then declare that the Pass is a *success* and stop. Otherwise, pick h' uniformly at random in $\mathcal{H} \setminus \{h\}$, and let v' be the vertex to which it is attached.

If $v' \notin \mathcal{F}$, then declare that a *collision* is found. Say that it is a *short collision* if moreover, $v' \in \overline{V}$; that it is a *long collision* otherwise. If either a short collision is found, or it is the second time during the Pass that a long collision is found, then declare that the Pass is a *failure*, do the updates displayed in (2.3) and stop. Else, go to Step 2.

If on the other hand $v' \in \mathcal{F}$, then do the following updates

$$\overline{V} \leftarrow \overline{V} \cup \{v'\}, \quad \overline{\mathcal{E}} \leftarrow \overline{\mathcal{E}} \cup \{h + h'\},$$

and go to Step 2.

Step 2

Do the following updates

$$\mathcal{H} \leftarrow \mathcal{H} \setminus \{h, h'\}, \quad \overline{\mathcal{H}} \leftarrow \overline{\mathcal{H}} \setminus \{h\}, \quad \mathcal{F} \leftarrow \mathcal{F} \setminus \{v'\},$$

and go to **Step 1** (it never happens that $h' \in \overline{\mathcal{H}}$ at this step, since we forbid short collisions, so there is no need to remove it from $\overline{\mathcal{H}}$).

We then iterate this pass, always starting with a new active seed, until every seed is inactive or has already been used as the starting point of a pass.

We now list some important observations concerning the effect of passes.

Observation 2.8 (A successful pass gives a favourable graph). Write $B_r(\rho)$ for the r -neighbourhood of ρ (in g). If the pass is a success, then the graph obtained by taking the union of $B_r(\rho)$ (seen as a sub-graph of g) and $(\overline{V}, \overline{\mathcal{E}})$, rooted at ρ , is (ℓ, r) -favourable. Indeed, this graph satisfies that

- x is in the vertex set,
- the distance from x to ρ is r ,
- the rooted graph $(x, (\overline{V}, \overline{\mathcal{E}}))$ is isomorphic to a rooted pruned ℓ -tree.

Moreover, such an (ℓ, r) -favourable graph does not intersect any possible (ℓ, r) -favourable graph obtained from previous passes.

Observation 2.9 (Many fresh sites). During one pass, at most

$$c_\ell := d + d^2 + \dots + d^\ell$$

iterations of step 1 are performed, so in particular at most c_ℓ sites are removed from the set \mathcal{F} of fresh sites. Initially, there are

$$n - |\mathcal{W}| [1 + (d+1) + \dots + (d+1)d^{r-1}]$$

fresh vertices. Since we start with $|\mathcal{W}|$ active seed vertices, we can run the pass no more than $|\mathcal{W}|$ times. At any given time, there are thus always at least $n - c_{r,\ell}|\mathcal{W}|$ fresh vertices, where

$$c_{r,\ell} := 1 + (d+1) + \dots + (d+1)d^{r-1} + c_\ell.$$

Observation 2.10 (Many passes). Apart from the bud x that is being explored, at most 2 active buds can be turned quiet during a pass. Indeed, a bud other than x can be turned quiet during a pass only when a collision occurs, since none of the buds ever belongs to \mathcal{F} . A collision can turn only one bud quiet at a time, and we allow for at most 2 collisions per pass. As a consequence, if we start (as we do) in a configuration where every bud is active and run the pass t times, then at most

$$t + \frac{2t}{(d+1)d^{r-1}}$$

seeds can be turned quiet, since initially every seed has $(d+1)d^{r-1}$ active buds. In other words, since we start with $|\mathcal{W}|$ seeds with all associated buds active, we can run the pass at least $\gamma_r|\mathcal{W}|$ times without running out of active seeds, where

$$(2.4) \quad \gamma_r = \left(1 + \frac{2}{(d+1)d^{r-1}}\right)^{-1}.$$

The crucial point for our reasoning is that we can take γ_r arbitrarily close to 1 by choosing r sufficiently large.

We now show that the probability to find many short collisions is very small. At any given time, there are no more than d^ℓ half-edges in $\overline{\mathcal{H}}$. The probability to find a short collision while running step 1 of the pass is thus bounded by

$$\frac{d^\ell}{(d+1)(n - c_{r,\ell}|\mathcal{W}|)} = \frac{d^\ell}{(d+1)(1 - kc_{r,\ell}\varepsilon)n},$$

using Observation 2.9 on the number of fresh sites.

As said in Observation 2.9, we run step 1 of the pass at most c_ℓ times, so the probability to find a short collision during one pass is bounded by

$$\mathbb{P} \left[\text{Bin} \left(c_\ell, \frac{d^\ell}{(d+1)(1 - kc_{r,\ell}\varepsilon)n} \right) \geq 1 \right].$$

As long as we choose ε sufficiently small that $kc_{r,\ell}\varepsilon < 1$, we see that this probability is bounded by \tilde{c}/n for some \tilde{c} depending on k, r and ℓ . Since we run at most $|\mathcal{W}|$ passes, the probability to find at least εn short collisions in total is bounded by

$$\mathbb{P} [\text{Bin} (|\mathcal{W}|, \tilde{c}/n) \geq \varepsilon n].$$

We then use Lemma 2.1 with $m = |\mathcal{W}| = k\varepsilon n$, $p = \tilde{c}/n$ and $p + \delta = 1/k$ to obtain that the probability above is bounded by

$$\exp \left(-|\mathcal{W}| \left[\frac{1}{k} \log \left(\frac{n}{k\tilde{c}} \right) + \left(1 - \frac{1}{k} \right) \log \left(\frac{1 - 1/k}{1 - \tilde{c}/n} \right) \right] \right),$$

which is much smaller than any given $(c\varepsilon^\eta)^{|\mathcal{W}|}$ when n is sufficiently large.

Let us say that a *double collision* occurs when two long collisions are found during one pass. We now turn to the estimation of the probability that there are many double collisions. It follows from Observation 2.9 that the probability to find a long collision while running step 1 of the pass is bounded by

$$\frac{(d+1)c_{r,\ell}|\mathcal{W}|}{(d+1)(n - c_{r,\ell}|\mathcal{W}|)} = \frac{c_{r,\ell}|\mathcal{W}|}{(n - c_{r,\ell}|\mathcal{W}|)},$$

Hence, the probability that a double collision is found during a pass is smaller than

$$\mathbb{P} \left[\text{Bin} \left(c_\ell, \frac{c_{r,\ell}|\mathcal{W}|}{(n - c_{r,\ell}|\mathcal{W}|)} \right) \geq 2 \right],$$

where we used the fact that during one pass, we run step 1 at most c_ℓ times. Recall that $|\mathcal{W}| = k\varepsilon n$, and note that

$$\mathbb{P} \left[\text{Bin} \left(c_\ell, \frac{k\varepsilon c_{r,\ell}}{1 - k\varepsilon c_{r,\ell}} \right) \geq 2 \right] \sim \frac{c'}{2}\varepsilon^2$$

as ε tends to 0, where

$$c' = (kc_\ell c_{r,\ell})^2.$$

Hence, we can choose ε sufficiently small (in terms of k, r and ℓ) so that

(2.5) the probability for a pass to find a double collision is bounded by $c'\varepsilon^2$.

The probability that we meet at least $\gamma_r|\mathcal{W}| - 2\varepsilon n$ double collisions in total is thus bounded by

$$\mathbb{P} [\text{Bin} (|\mathcal{W}|, c'\varepsilon^2) \geq \gamma_r|\mathcal{W}| - 2\varepsilon n] = \mathbb{P} [\text{Bin} (k\varepsilon n, c'\varepsilon^2) \geq (k\gamma_r - 2)\varepsilon n],$$

using (2.5) and the fact that we run at most $|\mathcal{W}|$ passes (and the fact that $|\mathcal{W}| = k\varepsilon n$ for the equality). Applying Lemma 2.1 with $m = k\varepsilon n = |\mathcal{W}|$, $p = c'\varepsilon^2$ and

$$p + \delta = \gamma_r - \frac{2}{k},$$

we see that this probability is bounded by

$$\exp\left(-|\mathcal{W}| \left[\left(\gamma_r - \frac{2}{k}\right) \log\left(\frac{\gamma_r - 2/k}{c'\varepsilon^2}\right) + \left(1 - \gamma_r + \frac{2}{k}\right) \log\left(\frac{1 - \gamma_r + 2/k}{1 - c'\varepsilon^2}\right) \right]\right).$$

By choosing r sufficiently large, we can have γ_r as close to 1 as we wish. In particular, we can make sure that

$$2\left(\gamma_r - \frac{2}{k}\right) > \eta,$$

so that the probability above is bounded by $(c\varepsilon^\eta)^{|\mathcal{W}|}$ for some constant c independent of ε .

To sum up, we have shown that in our procedure, we find

- at most εn short collisions and
- at most $\gamma_r |\mathcal{W}| - 2\varepsilon n$ double collisions

with probability larger than $1 - 2(c\varepsilon^\eta)^{|\mathcal{W}|}$. On this event, since we run the pass at least $\gamma_r |\mathcal{W}|$ times by Observation 2.10, we see that at least εn passes have been successful. This is exactly what we need to conclude, as was explained in Observation 2.8. \square

Corollary 2.11. *For every $k > 4$, every r sufficiently large and every ℓ , there exists $\varepsilon_0 > 0$ such that for every $\varepsilon \leq \varepsilon_0$, the following holds with \mathbb{P}_n -probability tending to 1 as n tends to infinity. From every r -prepared set $W \subseteq V_n$ of size $k\varepsilon n$, one can extract an (ℓ, r) -good subset of size εn .*

Proof. Proposition 2.7 ensures that there exists $\chi > 1$ such that for every r sufficiently large and every ℓ , there exists $c > 0$ such that for every $\varepsilon > 0$ sufficiently small, the probability for a given r -prepared set $W \subseteq V_n$ of size $k\varepsilon n$ not to contain any (ℓ, r) -good subset of size εn is smaller than $(c\varepsilon^\chi)^{k\varepsilon n}$. We prove the corollary by a union bound, counting the total number of sets $W \subseteq V_n$ of size $k\varepsilon n$. This number is

$$\binom{n}{k\varepsilon n} \leq \frac{n^{k\varepsilon n}}{(k\varepsilon n)!}.$$

Using the fact that $\log n! \geq n \log n - n$ (which can be proved by induction on n), we see that

$$\log \binom{n}{k\varepsilon n} \leq k\varepsilon n \log n - k\varepsilon n \log(k\varepsilon n) + k\varepsilon n = k\varepsilon n \left(1 + \log\left(\frac{1}{k\varepsilon}\right)\right).$$

If we choose $\varepsilon > 0$ sufficiently small, we can ensure that

$$(c\varepsilon^\chi)^{k\varepsilon n} \exp\left[k\varepsilon n \left(1 + \log\left(\frac{1}{k\varepsilon}\right)\right)\right]$$

tends to 0 as n tends to infinity (since $\chi > 1$), and this proves the result. \square

Proof of Proposition 2.6. Note first that the number of loops of size bounded by r remains tight as n tends to infinity (in order to see this, it suffices to check that the expectation of the number of such loops remains bounded). In particular, with probability tending to 1, there are less than \sqrt{n} loops of size bounded by r . We assume from now on that the event that there are no more than \sqrt{n} loops of size bounded by r is realized.

For any site $v \in V$, the size of its r -neighbourhood is bounded by

$$1 + (d+1) + (d+1)d + \dots + (d+1)d^{r-1} =: \bar{c}_r.$$

We choose $k > 8\bar{c}_r$, and observe that from any set of $k\varepsilon n$ vertices, we can extract a subset of $k\varepsilon n/\bar{c}_r$ vertices whose r -neighbourhoods are pairwise disjoint. Letting $k' = k/2\bar{c}_r > 4$ and using the fact that there are no more than \sqrt{n} loops of size bounded by r , we see that for n large enough, we can extract from any set of $k\varepsilon n$

vertices an r -prepared subset of cardinality $k'\varepsilon n$. The conclusion of the theorem then follows by Corollary 2.11. \square

Proof of Theorem 2.4. Let (o, H) be a rooted pruned ℓ -tree. Recall that H was obtained from \mathbb{T}_ℓ by removing an edge and taking the connected component of the root. Out of the d sub-trees of o in \mathbb{T}_ℓ , only one is affected by this construction. In particular, there exists (v, H') a rooted sub-tree of H such that v is a neighbour of o and (v, H') is isomorphic to $(o, \mathbb{T}_{\ell-1})$.

As a consequence, any (ℓ, r) -good set of vertices is $(\ell - 1, r + 1)$ -regenerative. Theorem 2.4 thus readily follows from Proposition 2.6. \square

3. SUBCRITICAL REGIME

We start from the following consequence of well-known estimates for the subcritical contact process on trees.

Proposition 3.1. *For any $\lambda < \lambda_1(\mathbb{T})$, there exists $C > 0$ such that*

$$\lim_{n \rightarrow \infty} \sup_{A \subseteq \mathbb{T}: |A|=n} P_{\mathbb{T}, \lambda} [\tau_{\mathbb{T}}^A > C \log n] = 0.$$

Proof. The proof will depend on the fact that, for any $\lambda < \lambda_1(\mathbb{T})$, there exist $c_0, C_0 > 0$ such that

$$(3.1) \quad \mathbb{E}_{\mathbb{T}, \lambda} [|\xi_t^o|] \leq C_0 e^{-c_0 t}, \quad t \geq 0.$$

This claim follows from putting together several results in Section I.4 of [Li99]; here we will simply outline them. We need the functions $\phi_\lambda(\rho)$ and $\beta(\lambda)$ (where $\rho > 0$ is an extra variable which is not involved in the construction of the contact process); these functions are defined respectively in equation (4.23), page 87, and (4.48), page 96. In Proposition 4.27(a) and (b) we respectively have

$$(3.2) \quad \phi_\lambda(\rho) = \phi_\lambda \left(\frac{1}{\rho d} \right)$$

and

$$(3.3) \quad \mathbb{E}_{\mathbb{T}, \lambda} [|\xi_t^o|] \leq C_\lambda \cdot \phi_\lambda(1)^t, \quad t \geq 0,$$

for some constant $C_\lambda > 0$. In Proposition 4.44(a), we have

$$(3.4) \quad \rho < \rho', \quad \phi_\lambda(\rho') \geq 1 \implies \phi_\lambda(\rho) < \phi_\lambda(\rho').$$

Regarding β , in Corollary 4.78, Theorem 4.83 and Theorem 4.130, we respectively have that

$$(3.5) \quad \beta(\lambda_1) = \frac{1}{d},$$

$$(3.6) \quad \phi_\lambda(\beta(\lambda)) = \phi_\lambda \left(\frac{1}{\beta(\lambda) \cdot d} \right) = 1,$$

$$(3.7) \quad \beta \text{ is strictly increasing in } [0, \lambda_1].$$

Now, for $\lambda < \lambda_1$, (3.5) and (3.7) give $\beta(\lambda) < \frac{1}{d}$. Together with (3.6), this shows that $\phi_\lambda(\rho) = 1$ for some $\rho > 1$. Then, using (3.4), we get $\phi_\lambda(1) < 1$, and then (3.3) gives the desired equation (3.1).

Let us now show how (3.1) completes the proof of our proposition. Noting that $|\xi_t^A| = |\cup_{x \in A} \xi_t^x| \leq \sum_{x \in A} |\xi_t^x|$, we have

$$P_{\mathbb{T}, \lambda} [|\xi_t^A| \neq \emptyset] \leq E_{\mathbb{T}, \lambda} [|\xi_t^A|] \leq |A| \cdot M_t \leq C_0 |A| e^{-c_0 t}.$$

The proof is completed by taking $C = 2/c_0$. \square

Part (a) of Theorem 1.2 is a consequence of the above proposition and the following result:

Lemma 3.2. *For any finite graph $G = (V, E)$ with degree bounded by $d + 1$, $A \subseteq V$ and $t > 0$,*

$$P_{G,\lambda} [\tau_G^A > t] \leq \sup_{B \subseteq T: |B|=|A|} P_{T,\lambda} [\tau_T^B > t].$$

Proof. Since the contact process is unaffected by the presence of loops (edges that start and end at the same vertex), we assume that G has none.

We will now recall the concept of universal covering of the graph G ; we will construct from G a new graph $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with certain desirable properties. We start fixing a reference vertex $x \in V$. We say that a sequence $\gamma = (\vec{e}_1, \dots, \vec{e}_n)$ of oriented edges of \vec{E} is a *non-backtracking path* from x if $v_0(\vec{e}_1) = x$ and, for $1 \leq i < n$, $v_1(\vec{e}_i) = v_0(\vec{e}_{i+1})$ and $u(\vec{e}_i) \neq u(\vec{e}_{i+1})$ (recall that, for an oriented edge \vec{e} , $v_0(\vec{e})$, $v_1(\vec{e})$ and $u(\vec{e})$ respectively denote the starting vertex, ending vertex and undirected edge associated to \vec{e}). Let \mathcal{V} be the set of all non-backtracking paths from x , including an empty path which we denote by o . For any $\gamma, \gamma' \in \mathcal{V}$ with $\gamma = (\vec{e}_1, \dots, \vec{e}_n)$ and $\gamma' = (\vec{e}_1, \dots, \vec{e}_n, \vec{e}_{n+1})$, we connect γ and γ' by an edge; this defines the edge set \mathcal{E} of \mathcal{T} . Finally, put $\psi(o) = x$ and, for $\gamma = (\vec{e}_1, \dots, \vec{e}_n)$, put $\psi(\gamma) = v_1(\vec{e}_n)$, the ending vertex of the path γ . It is now easy to check that \mathcal{T} and ψ satisfy the properties:

- (a) \mathcal{T} is a tree with degree bounded by $d + 1$;
- (b) for every $\gamma \in \mathcal{V}$, ψ maps the neighbourhood of γ bijectively to the set

$$\{\vec{e} : \vec{e} \in \vec{E} \text{ has } \psi(\gamma) \text{ as its starting vertex}\}$$

(in case $\psi(\gamma)$ is connected to each of its neighbours by a single edge, property (b) just says that ψ maps the neighbourhood of γ bijectively to the neighbourhood of $\psi(\gamma)$).

For $y \in V$, the set $\psi^{-1}(y)$ is called the *fiber* of y . Define the set of configurations of $\{0, 1\}^{\mathcal{V}}$ that have at most one particle per fiber,

$$\Omega_{\mathcal{T}} = \left\{ \zeta \in \{0, 1\}^{\mathcal{V}} : \sum_{\gamma \in \psi^{-1}(v)} \zeta(\gamma) \in \{0, 1\} \right\}.$$

Define the projection $\pi : \Omega_{\mathcal{T}} \rightarrow \{0, 1\}^V$ by $[\pi(\zeta)](v) = \sum_{\gamma \in \psi^{-1}(v)} \zeta(\gamma)$ for $v \in V$. We abuse notation and, for a set $B \subseteq \mathcal{V}$, we write $B \in \Omega_{\mathcal{T}}$ if $\mathbf{1}_B \in \Omega_{\mathcal{T}}$.

Given $B \in \Omega_{\mathcal{T}}$, we define the *constrained contact process on \mathcal{T}* , $(\zeta_t^B)_{t \geq 0}$, as follows. We set $\zeta_0^B = B$ and let ζ evolve as a contact process on \mathcal{T} with the restriction that we suppress every transition which would result in a configuration not in $\Omega_{\mathcal{T}}$, that is, births on vertices belonging to fibers already containing infected vertices. Formally, (ζ_t) has generator

$$\mathcal{L}f(\zeta) = \sum_{\gamma \in \mathcal{V}} [f(\zeta^{0 \rightarrow \gamma}) - f(\zeta)] + \lambda \cdot \sum_{\substack{\{\gamma, \gamma'\} \in \mathcal{E}: \\ \zeta(\gamma) = 1}} \left[\mathbf{1}_{\{\zeta^{1 \rightarrow \gamma'} \in \Omega_{\mathcal{T}}\}} \cdot (f(\zeta^{1 \rightarrow \gamma'}) - f(\zeta)) \right],$$

where $\zeta^{i \rightarrow \gamma'}$ is the configuration obtained by modifying ζ so that γ' is set to state i . Noting that G is finite, and hence ζ has at most $|V|$ infected vertices at any given time, it is not hard to see that the above generator indeed gives rise to a Feller process on $[0, \infty)^{\Omega_{\mathcal{T}}}$. In fact, ζ can be constructed from a Harris system on \mathcal{T} , with the extra care of ignoring any transmission mark which would cause a fiber to become doubly occupied. Also, using property (b) of \mathcal{T} and ψ stated above, it is easy to show that $(\pi(\zeta_t^B))_{t \geq 0}$ has the same distribution as the contact process on G started from $\pi(B)$.

Now, if we start from $A \subseteq V$, we can choose an arbitrary $B \in \Omega_{\mathcal{T}}$ such that $\pi(B) = A$ and conclude that τ_G^A has the same distribution as $\inf\{t : \zeta_t^B = \emptyset\}$. By

seeing \mathcal{T} (and hence B) as a subset of \mathbb{T} , we have that ζ^B is stochastically dominated by the contact process on \mathbb{T} started from B infected, and hence $\inf\{t : \zeta_t^B = \emptyset\}$ is stochastically dominated by $\tau_{\mathbb{T}}^B$. This completes the proof. \square

Let G be a graph and x a vertex of G . Given the contact process on G started from a single infection at x , $(\xi_t^x)_{t \geq 0}$, define

$$\kappa_G^x = \sup\{\text{dist}(x, y) : y \in \xi_t^x \text{ for some } t \geq 0\}.$$

Lemma 3.3. *For any finite graph G with degree bounded by $d + 1$, any vertex x of G and $k > 0$,*

$$\mathbb{P}[\kappa_G^x > k] \leq \mathbb{P}[\kappa_{\mathbb{T}}^o > k].$$

Proof. Repeating the construction in the previous lemma, we note that, if $\psi(\gamma) = x$ and $\psi(\gamma') = y$, then $\text{dist}_{\mathcal{T}}(\gamma, \gamma') \geq \text{dist}_G(x, y)$, and we then see that κ_G^x is stochastically dominated by $\sup\{\text{dist}_{\mathcal{T}}(\gamma, \gamma') : \gamma' \in \zeta_t^{\{\gamma\}} \text{ for some } t \geq 0\}$. The latter is in turn stochastically dominated by $\kappa_{\mathbb{T}}^o$, completing the proof. \square

Proof of Theorem 1.4. In order to show that the contact process on G dies out, it suffices to show, for any vertex x , that $\mathbb{P}[\tau_G^x < \infty] = 1$. Denote by $B_G(x, r)$ the subgraph of G induced by the set of vertices at graph distance less than r from x . Then,

$$\mathbb{P}[\tau_G^x < \infty] = \lim_{r \rightarrow \infty} \mathbb{P}[\kappa_G^x \leq r] = \lim_{r \rightarrow \infty} \mathbb{P}[\kappa_{B_G(x, 2r)}^x \leq r] \geq \lim_{r \rightarrow \infty} \mathbb{P}[\kappa_{\mathbb{T}}^x \leq r] = 1.$$

\square

REFERENCES

- [Be] I. Benjamini. *Coarse geometry and randomness – École d’Été de Probabilités de Saint-Flour 2011*. Lecture Notes in Mathematics **2100**, Springer (2013).
- [BNP11] I. Benjamini, A. Nachmias, Y. Peres. *Is the critical percolation probability local?* Probab. Theory Related Fields **149** (1-2), 261–269 (2011).
- [BS01] I. Benjamini, O. Schramm. *Recurrence of distributional limits of finite planar graphs*. Electron. J. Probab. **6** (23), 13 pp. (2001).
- [CGOV84] M. Cassandro, A. Galves, E. Olivieri, M. Vares. *Metastable behavior of stochastic dynamics: a pathwise approach*. J. Statist. Phys. **35** (5-6), 603-634 (1984).
- [CD09] S. Chatterjee, R. Durrett. *Contact processes on random graphs with power law degree distributions have critical value 0*. Ann. Probab. **37**, 2332 - 2356 (2009).
- [Ch94] J.W. Chen. *The contact process on a finite system in higher dimensions*. Chinese J. Contemp. Math. **15**, 13-20 (1994).
- [CMMV13] M. Cranston, T. Mountford, J.C. Mourrat, D. Valesin, *The contact process on finite trees revisited*. Preprint, available at <http://www.math.ubc.ca/~valesin/cptreerev.pdf>
- [DL88] R. Durrett, X. Liu. *The contact process on a finite set*. Ann. Probab. **16**, 1158-1173 (1988).
- [DS88] R. Durrett, R.H. Schonmann. *The contact process on a finite set II*. Ann. Probab. **16**, 1570-1583 (1988).
- [Li99] T. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*. Grundlehren der mathematischen Wissenschaften **324**, Springer (1999).
- [MT13] S. Martineau, V. Tassion. *Locality of percolation for abelian Cayley graphs*. Preprint, arXiv:1312.1946 (2013).
- [Mo93] T. Mountford. *A metastable result for the finite multidimensional contact process*. Canad. Math. Bull. **36** (2), 216-226 (1993).
- [Mo99] T. Mountford. *Existence of a constant for finite system extinction*. J. Statist. Phys. **96** (5-6), 1331-1341 (1999).
- [MMVY12] T. Mountford, J.C. Mourrat, D. Valesin, Q. Yao. *Exponential extinction time of the contact process on finite graphs*. Preprint: arXiv:1203.2972 (2012).
- [MVY13] T. Mountford, D. Valesin, Q. Yao. *Metastable densities for the contact process on power law random graphs*. Electronic Journal of Probability, Vol 18 1-36 (2013).
- [SS98] M. Salzano, R. Schonmann. *A new proof that for the contact process on homogeneous trees, local survival implies complete convergence*. Ann. Probab. **26**, 1251-1258 (1998).
- [Sc85] R. Schonmann. *Metastability for the contact process*. J. Statist. Phys. **41** (3-4), 445-464 (1985).

- [St01] A. Stacey. *The contact process on finite homogeneous trees*. Probab. Theory Related Fields. **121** (4), 551-576 (2001).

(Jean-Christophe Mourrat) ENS LYON, CNRS, 46 ALLÉE D'ITALIE, 69007 LYON, FRANCE

(Daniel Valesin) UNIVERSITY OF BRITISH COLUMBIA, 1984 MATHEMATICS ROAD, VANCOUVER BC, CANADA V6T 1Z2